

Modeling Short Over-Dispersed Spike Count Data: A Hierarchical Parametric Empirical Bayes Framework

Qi She, Beth Jelfs, Rosa H.M.Chan

September 27, 2016

Abstract

In this letter, a Hierarchical Parametric Empirical Bayes model is proposed to model spike count data. We have integrated Generalized Linear Models (GLMs) and empirical Bayes theory to simultaneously provide three advantages: (1) a model of over-dispersion of spike count values; (2) reduced MSE in estimation when compared to using the maximum likelihood method for GLMs; and (3) an efficient alternative to inference with fully Bayes estimators. We apply the model to study both simulated data and experimental neural data from the retina. The simulation results indicate that the new model can estimate both the weights of connections among neural populations and the output firing rates (mean spike count) efficiently and accurately. The results from the retinal datasets show that the proposed model outperforms both standard Poisson and Negative Binomial GLMs in terms of the prediction log-likelihood of held-out datasets.

1 Introduction

The functionality of a neural network depends on its connectivity, hence, if we want to be able to investigate different properties of the network it is important to be able to accurately reconstruct the connectivity. Functional connectivity provides a strategy to construct the connectivity, which rather than focusing on the anatomical connectivity, takes an input-output driven approach. In this

way, the connectivity is analyzed, directly from the neural activity, from the point of view of the statistical dependences among the data [1, 2]. Crucial to the success of such an approach is how to model the relationship between the spike count data in order to build a reasonable estimate of these dependences.

Of the various approaches to modeling the statistical dependences between multiple spike-train data, all require certain assumptions about the relationships within the data. However, with the recent increase in accessibility of datasets containing simultaneously recorded spike-train data, it is now possible to reconsider, and test, the effectiveness of different novel methods for extracting functional dependences at the neuronal level. Two of the most commonly used methods are Generalized Linear Models (GLMs) and Latent Variable Models (LVMs). GLMs have the benefit of being able to take into account stimuli, self-spiking history, and outputs of the trigger neurons [3, 4, 5, 6, 7]. While the LVMs can also include self-spiking history and external stimuli, it assumes that multiple neural activities arise from some common hidden sources. Accordingly, this method can provide us with a low dimensional representation of observed spiking activities. Nonetheless, inferring the posterior distribution of the latent variables and estimating model parameters can be difficult. In the literature there has been a great deal of discussion as to how best to accurately and efficiently infer and learn the LVMs [8, 9, 10, 11].

In computational neuroscience it is currently common practice, when using both the GLMs and LVMs, to assume that the spike count data at each time stamp are Poisson distributed. While the Poisson assumption works well in the majority of cases [12], the Poisson model is established on the assumption that at each time stamp the mean spike count is equal to its variance. This ignores the fact that in some cases the variance of the spike count could be much larger than its mean [13], that is, the data is over-dispersed. As an alternative, the Negative Binomial (NB) model has been proposed as a solution to handling over-dispersion cases when modeling spike count data [14].

Despite the advantages of the NB model, when the spike-train data is sparse, and a large number of neurons are recorded simultaneously, the accuracy of the estimated coefficients using GLMs with NB response is low [15]. At the same time, under GLMs the estimator for the parameter of the NB distribution is derived entirely based on the sample means. Hence, if the length of the data is limited, be that either due to the length of the experiment or the need for real-time inference [16, 17, 18, 19], the estimator of the mean spike count can have a large mean square error (MSE).

Therefore, when the model selected does not match the simulated data, further modifications need to be made to the model. One approach to address such modifications is to employ shrinkage, that is the use of other information to ‘shrink’ the distance of the estimate towards the information, as a means to optimize the bias-variance trade-off. If we have subjective information about our parameters, using a Bayesian model can provide advantageous prior information.

Both fully and empirical Bayes inference produce shrinkage estimators which combine observations with prior information to achieve better estimation. Fully Bayes inference assumes the hyperparameters of the prior distribution have hyperprior distributions, which need to be integrated out; while empirical Bayes sets the parameters in the highest level of the hierarchical model with their most likely value. Thus, as often we cannot obtain the closed form of the posterior distribution, fully Bayes inference requires a sampling strategy to approximate this distribution. Correspondingly, this comes at a high computational cost, especially for high-dimensional data [20]. Conversely, in empirical Bayes the hyperparameters are obtained by maximizing marginal likelihood function with much lower computational cost. Hence, by combining empirical Bayes with the NB GLM we can produce an estimator for the probability parameter, θ_i , of the NB distribution which should efficiently handle both over-dispersion and short data lengths.

In this letter, we propose a hierarchical empirical Bayes estimator for **Short Over-Dispersed Spike Counts**, which we call the “**SODS**” estimator. The hierarchical framework gives a prior distribution on the probability parameter (θ_i) of the NB distribution, and using empirical Bayes inference provides a shrinkage estimator of θ_i . The hyperparameters of the prior distribution are estimated using maximum marginal likelihood methods. The estimated value of θ_i can then be used to obtain the mean spike count. The purpose of this approach is to achieve three goals:

1. Modeling of over-dispersed data where the spike count variance is greater than its mean value;
2. Reducing the MSE of GLMs when the data length is short;
3. Improving the efficiency of estimation and inference for both simulated and experimental datasets.

This letter is organised as follows. In Section 2 we introduce the **SODS**. Section 3 discusses estimation of the unknown parameters of the **SODS**, via numerical optimization of the maximum




Figure 1: Network Model and Graphical representation. *left panel*: Neuron #a is the target neuron, Neuron #b, #c, #d and #e are trigger neurons used as input regressors. The goal of the model is to estimate the underlying firing probability of #a, and derive the corresponding spike count. The observed data are multiple spike-train data recorded simultaneously; *right panel*: Graphical visualization of the proposed model.

marginal likelihood, and the roles of these variables. Results for both simulated and experimental data are presented in Section 4. Discussion of the main contributions and findings of this letter are given in Section 5. Finally Section 6 concludes the letter.

2 Modeling

In order to derive the **SODS** for estimating probability parameter of NB distribution we first introduce the **SODS** model and then discuss the advantages of this model and the performance of our empirical Bayes estimator. The left panel of Fig. 1 illustrates the network considered in this letter in its simplest case¹, for the purposes of our estimator we consider the connections among the neural populations denoted as “functional dependences”. Based on this representation the aim of the **SODS** model is to estimate the underlying firing probability and the mean spike count of the target neuron. Our goal is to find the statistical dependences between multiple spike-train data, and use other neurons’ spiking activities as regressors to estimate the target neuron’s firing rate (mean spike count) at each time stamp.

¹This 5-neuron case is a simple one used only for visualization. In the simulations, 10, 20, or 50 neurons are involved in the networks

The right panel of Fig. 1 shows the hierarchical structure of our model, and the key variables and parameters of the “functional dependences”. Table 1 provides a complete summary of all the variables used in the **SODS** estimator. The data length (number of time bins) is equal to K , at the first level of the hierarchical model μ_i , the mean of the Beta prior of the NB probability parameter (θ_i) is formed from the GLM of the input regressors x_i , the weight vector ω , and the link function determined by γ . At the next level the actual firing probability, $(1 - \theta_i)$, is controlled by σ (the degree of freedom of the prior distribution) and prior distribution (Beta distribution). Finally, the probability parameter θ_i , together with the key parameter for the NB distribution r , generate the observed spike count data Y_i .

2.1 Conjugate Beta Prior for Negative Binomial Data

cite, several different works have tried to model non-Poisson spike counts [21, 22, 23, 24]. In particular, the NB response has been implemented to model spike count data in the visual pathway [25], which is beneficial for understanding the representation of sensory information in the brain. If we consider the case when the spike count data is over-dispersed, then modeling the neuronal response generated from the NB distribution is a natural choice, as the NB distribution assumes that the variance of the spike count data is larger than its mean value. At the same time the NB distribution can also be considered to be a generalization of the Poisson distribution where the rate parameter of the Poisson distribution is represented as a Gamma distribution. Thus, the Gamma-Poisson mixture, models the spike count of the j -th trial at time i , denoted as y_{ij} , as a Poisson distribution. Consequently, the mean value, λ_i , of the count data from time bin i , becomes a Gamma distribution. These processes are given as

$$y_{ij} \mid \lambda_i \sim \text{Pois}(\lambda_i) = \frac{\lambda_i^{y_{ij}}}{y_{ij}!} e^{-\lambda_i}, \quad (1)$$

$$\lambda_i \mid r, \theta_i \sim \text{Ga}\left(r, \frac{\theta_i}{1 - \theta_i}\right) = \lambda_i^{r-1} \frac{\exp\left(-\lambda_i \frac{\theta_i}{1 - \theta_i}\right)}{\left(\frac{1 - \theta_i}{\theta_i}\right)^r \Gamma(r)}. \quad (2)$$

By combining these two processes, we can model the spike count response as a NB distribution in the form

$$y_{ij} \mid r, \theta_i \sim \text{NB}(r, \theta_i) = \binom{r + y_{ij} - 1}{y_{ij}} (\theta_i)^r (1 - \theta_i)^{y_{ij}}, \quad (3)$$

Table 1: Summary of variable definitions.

Variable	Definition
i	Time bin
j	Trial number at each time bin
y_{ij}	Spike count data of j-th trial at i-th time bin
\mathbf{x}_i	Vector of regressors at i-th time bin
λ_i	Mean of Poisson distribution
$1 - \theta_i$	Firing probability; θ_i : probability parameter of NB distribution
r	The number of failures in NB distribution
α_i, β_i	Parameters of Beta distribution
σ	Degree of freedom of prior distribution ($\sigma = \alpha_i + \beta_i$)
ω	Vector of weights
$g(\cdot)$	Family of link functions
γ	Parameter determining specific form of link function
μ_i	Mean of prior distribution
n_i	Number of trials at i-th time bin
\bar{y}_i	Mean spike count across all trials
π_i	Weight of the observation component of proposed estimator at i-th time bin
K	Data length (the total number of bins)
A_i, B_{ij}, C_{ij}	Components of the gradients
p	Element number of ω
N	Total number of neurons
ζ	$(r, \omega, \sigma, \gamma)$
\mathbf{k}	Vector of Lagrangian coefficients
m	Variable number from 1 to $N + 3$
t	Iteration step
\mathbf{d}	Vector of search direction
Δ_t	Step length of iteration

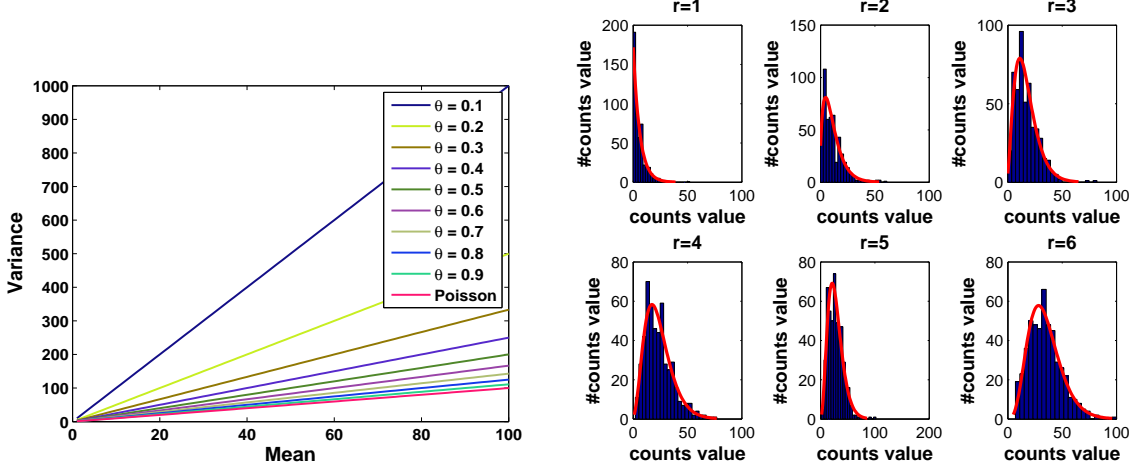


Figure 2: *left panel*: Relationship between the mean and variance for different settings of θ . Compared with the Poisson model the NB model shows super-Poisson variability (the variance is larger than the mean for the same θ). The smaller θ the more significant the super-Poisson effect; *right panel*: NB distribution with different r and the same θ .

where $(1 - \theta_i)$ represents the firing probability at time i , r is the number of failures and j is the trial number. The effect of decreasing the value of θ_i (with a fixed value of r) on the mean and variance of the NB distribution is shown in the left panel of Fig. 2, illustrating the super-Poisson variability which makes this model appropriate for use with over-dispersed data.

Our hierarchical model estimates the prior over the probability parameter θ_i (which controls the firing probability) directly from the data under an empirical Bayes framework. We have selected the Beta distribution as the prior distribution for θ_i as it is the conjugate prior of the NB distribution of Eq. (3). Hence, making derivation of a closed-form posterior distribution possible. The Beta distribution of θ_i is given by

$$\theta_i \sim \text{Beta}(\alpha_i, \beta_i) = \frac{\theta_i^{\alpha_i-1} (1 - \theta_i)^{\beta_i-1}}{B(\alpha_i, \beta_i)}, \quad (4)$$

$$B(\alpha_i, \beta_i) = \frac{\Gamma(\alpha_i)\Gamma(\beta_i)}{\Gamma(\alpha_i + \beta_i)}, \quad (5)$$

where α_i and β_i are the hyperparameters of the distribution. The Beta function, $B(\alpha_i, \beta_i)$, is equivalent to $B(\alpha_i, \beta_i) = \int_0^1 x^{\alpha_i-1} (1-x)^{\beta_i-1} dx$, and can also be represented using Gamma function $\Gamma(\cdot)$, as shown in Eq. (5), where $\Gamma(t) = \int_0^\infty x^{t-1} e^{-x} dx$.

In order to allow us to control the degrees of freedom of the prior distribution we introduce another hyperparameter, σ . The mean of the prior, $\mathbb{E}(\theta_i | \alpha_i, \beta_i) = \mu_i$, is given by $\mu_i = \frac{\alpha_i}{\alpha_i + \beta_i}$. If we assume this degree of freedom of the distribution can be given by $\sigma = \alpha_i + \beta_i$, then we can rearrange to get $\alpha_i = \sigma \mu_i$ and $\beta_i = \sigma(1 - \mu_i)$. Forming the prior in this manner, allows us to determine the form of the Beta prior distribution based purely on the mean of the prior, μ_i and the hyperparameter, σ .

In order to calculate the prior mean, we use a GLM as a predictor. The GLM is formed from a given vector of input regressors, \mathbf{x}_i , typically spike counts from other neurons, and a vector of weights, $\boldsymbol{\omega}$, which capture the directed effect of the neural population on the target neuron. As we are using the Beta distribution as the prior for θ_i , it is not only reasonable but also necessary to scale the underlying firing probability between 0 and 1 via a link function. However, we do not want to constrain the link function to be a fixed form, hence, we select a family of link functions governed by a hyperparameter, γ . Thus, the prior mean becomes

$$\mu_i = g^{-1}(\mathbf{x}_i^T \boldsymbol{\omega}) = 1 - \left(\gamma e^{\mathbf{x}_i^T \boldsymbol{\omega}} + 1 \right)^{-\frac{1}{\gamma}}. \quad (6)$$

The hyperparameter γ , is a flexible parameter which determines the specific form of the link function, $g(\cdot)$. We select this link family in particular as it can represent many widely used link functions, for instance, the logit function when $\gamma = 1$. It also constrains the prior mean, modeled as the mean value of the firing probability, to $\mu_i > 0$. The exact form of the hyperparameter, γ , is estimated directly from the data, therefore ensuring the flexibility of the nonlinear transformation of the regressors.

2.2 Empirical Bayes Estimator: SODS

Having used the GLM to form a predictor for μ_i , the next step in deriving our **SODS** estimator is to address the posterior distribution of θ_i . As a consequence of the Beta distribution being the conjugate prior for the NB likelihood function, we can derive the closed-form of the posterior distribution of θ_i given $Y_{ij} = y_{ij}$, in the form of the Beta Negative Binomial distribution

$$\theta_i | y_{ij} \sim \text{Beta} \left(\alpha_i + n_i r, \beta_i + \sum_{j=1}^{n_i} y_{ij} \right) = \text{Beta} \left(\sigma \mu_i + n_i r, \sigma (1 - \mu_i) + n_i \bar{y}_i \right), \quad (7)$$

where n_i is the number of trials of the i -th time bin and \bar{y}_i is the mean value of the count data across all trials at each bin, i . Substituting the inverse link function of μ_i into Eq. (7) we get

$$\theta_i \mid y_{ij}, r, \omega, \sigma, \gamma \sim \text{Beta} \left(\sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma) + n_i r, \sigma \left(1 - g^{-1}(\mathbf{x}_i^T \omega, \gamma) \right) + n_i \bar{y}_i \right). \quad (8)$$

Our proposed **SODS** estimator is defined as the expectation of the conditional posterior for θ_i . The **SODS** estimator is given as

$$\theta_i^{\text{SODS}} = \mathbb{E}(\theta_i \mid y_{ij}, r, \omega, \sigma, \gamma) = \frac{n_i r + \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)}{n_i r + n_i \bar{y}_i + \sigma}. \quad (9)$$

The estimator has a group of unknown parameters $\{r, \omega, \sigma, \gamma\}$ which need to be estimated, these are, respectively, the number of failures in NB distribution (r); the vector of weights in the GLM (ω); the degree of freedom of the prior distribution (σ); and the parameter determining the specific form of the link function in the GLM (γ). A more thorough description of their functions and effects on the estimator will be given in Section 3.2. After proposing the empirical Bayes estimator for the hyperparameters of θ_i , we can derive the mean spike count based on θ_i^{SODS} , according to Eq. (3) as

$$\mathbb{E}[Y_i \mid \theta_i^{\text{SODS}}] = \hat{r} \left(\frac{1}{\theta_i^{\text{SODS}}} - 1 \right) = \hat{r} \left(\frac{n_i \bar{y}_i + \hat{\sigma} - \hat{\sigma} g^{-1}(\mathbf{x}_i^T \hat{\omega}, \gamma)}{n_i \hat{r} + \hat{\sigma} g^{-1}(\mathbf{x}_i^T \hat{\omega}, \hat{\gamma})} \right). \quad (10)$$

Having obtained the final estimator for mean spike count in Eq. (10), we now address how the structure of the estimator θ_i^{SODS} can be used to explain its advantages. We can view the estimator as a combination of the prior mean GLM and the observed data. In order to have a clear representation of our proposed estimator, we rewrite Eq. (9) as a convex combination of the two components (the prior mean GLM and the spike count observations):

$$\begin{aligned} \theta_i^{\text{SODS}} &= \frac{n_i r + \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)}{n_i r + n_i \bar{y}_i + \sigma} \\ &= \left(\frac{n_i r + n_i \bar{y}_i}{n_i r + n_i \bar{y}_i + \sigma} \right) \left(\frac{n_i r}{n_i r + n_i \bar{y}_i} \right) + \left(\frac{\sigma}{n_i r + n_i \bar{y}_i + \sigma} \right) \left(g^{-1}(\mathbf{x}_i^T \omega, \gamma) \right). \end{aligned} \quad (11)$$

If we assign the weight of the observation component as $\pi_i = \frac{n_i r + n_i \bar{y}_i}{n_i r + n_i \bar{y}_i + \sigma}$, then $1 - \pi_i = \frac{\sigma}{n_i r + n_i \bar{y}_i + \sigma}$, and Eq. (11) can be simplified to give

$$\theta_i^{\text{SODS}} = \pi_i \left(\frac{n_i r}{n_i r + n_i \bar{y}_i} \right) + (1 - \pi_i) \left(g^{-1}(\mathbf{x}_i^T \omega, \gamma) \right). \quad (12)$$

Hence, the value π_i controls the weights assigned to the observations, $\frac{n_i r}{n_i r + n_i \bar{y}_i}$, and to the prior mean value, $g^{-1}(x_i^T \omega, \gamma)$ and $0 \leq \pi_i \leq 1$.

We can consider π_i as a coefficient of shrinkage and, thus, θ_i^{SODS} as a shrinkage estimator. When $\pi_i = 1$, it results in no shrinkage of θ_i^{SODS} from the observations to the prior mean value. The value of π_i is determined by the final estimated values of the degree of freedom of the prior distribution σ , and the number of failures r . That is to say, when $\sigma \gg n_i r + n_i \bar{y}_i$, the estimator is approaching the prior mean, while when $\sigma \ll n_i r + n_i \bar{y}_i$, the estimator is determined by the observations. This structure has the advantage that all the hyperparameters (determining the weights π_i) are learned from the dataset and the final estimator is obtained according to the weights π_i relating the prior information and the observations.

Furthermore, if we consider the conditional posterior variance of θ_i , we can describe this variance as

$$\text{Var}(\theta_i | y_{ij}, r, \omega, \sigma, \gamma) = \frac{\left(\frac{n_i r + \sigma g^{-1}(x_i^T \omega, \gamma)}{n_i r + n_i \bar{y}_i + \sigma} \right)^2 - \left(\frac{n_i r + \sigma g^{-1}(x_i^T \omega, \gamma)}{n_i r + n_i \bar{y}_i + \sigma} \right)^2}{n_i r + n_i \bar{y}_i + \sigma + 1}. \quad (13)$$

Then, assuming the estimations of r and σ are accurate, the value of $n_i r + n_i \bar{y}_i + \sigma$ is fixed. Thus, the conditional posterior variance of the firing probability, θ_i , becomes a quadratic function of only the prior mean value GLM. The variance achieves the global maximum when $g^{-1}(x_i^T \omega, \gamma) = \frac{1}{2} n_i \bar{y}_i - \frac{1}{2} n_i r + \frac{1}{2}$, and

$$\text{Var}_{\max}(\theta_i | y_{ij}, r, \omega, \sigma, \gamma) = \frac{1}{4 \times (n_i r + n_i \bar{y}_i + \sigma + 1)}. \quad (14)$$

3 Empirical Bayes Estimation

After deriving the **SODS** estimator from Section 2.2, we need to estimate the parameters $r, \omega, \gamma, \sigma$. In order to achieve this we use an empirical Bayes approach which allows us to estimate the hyperparameters directly from the recorded data. The first step is the estimation of the unknown parameters using maximum marginal likelihood, where the marginal distribution is the spike count response conditional on the hyperparameters. After that, we give the derivation of the gradient functions for each of the unknown parameters. Finally, we talk about the roles of each of the parameters of our model and how to use prior knowledge to set the initial value to give more stable and accurate optimization results.

The unknown parameters can be estimated by utilizing the marginal distribution, this has the benefit of having relatively low computational costs, as using the maximum marginal likelihood approach does not include any assumptions on the hyperparameters. The marginal density function $f_i(y_{ij}|r, \omega, \sigma, \gamma)$ is derived from the conjugate prior density from Eq. (4) and the NB density (likelihood) described in Eq. (3), with the marginal distribution in this density function being the Beta Negative Binomial distribution. Thus, the marginal density of the spike count conditional on the unknown parameters is given as

$$f_i(y_{ij}|r, \omega, \sigma, \gamma) = \frac{B\left(r + \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma), y_{ij} + \sigma - \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)\right)}{B(r, y_{ij})B\left(\sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma), \sigma - \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)\right)y_{ij}}, \quad (15)$$

where $B(\alpha_i, \beta_i) = \frac{\Gamma(\alpha_i)\Gamma(\beta_i)}{\Gamma(\alpha_i + \beta_i)}$. Conditioning on y_{ij} , we get the log marginal likelihood of the conditional posterior distribution as

$$\begin{aligned} \ell(r, \omega, \sigma, \gamma) \propto & \sum_{i=1}^K \sum_{j=1}^{n_i} \left[\log \Gamma\left(r + \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)\right) \right. \\ & + \log \Gamma\left(y_{ij} + \sigma - \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)\right) + \log \Gamma(r + y_{ij}) \\ & + \log \Gamma(\sigma) - \log \Gamma(r + y_{ij} + \sigma) - \log \Gamma(r) \\ & \left. - \log \Gamma\left(\sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)\right) - \log \Gamma\left(\sigma - \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)\right) \right], \quad (16) \end{aligned}$$

where K is the data length (the number of time bins), and n_i is the number of trials at time i . We estimate the parameters $r, \omega, \sigma, \gamma$, by maximizing the log marginal likelihood, as shown in Eq. (16), and then substitute them into the **SODS** estimator.

3.1 Quasi-Newton Numerical Optimization

As previously stated, estimation of the hyperparameters $\hat{r}, \hat{\omega}, \hat{\sigma}, \hat{\gamma}$ is conducted by maximizing the marginal likelihood function. Normally, in hierarchical modeling, it is hard to obtain closed-forms for the estimators of the hyperparameters. Thus, numerical optimization is applied to search for optimal hyperparameters. Here, we need to approximate the Hessian matrix at each iteration, which is achieved by applying the Quasi-Newton optimization method to the approximation of the Hessian matrix. The *Davidon-Fletcher-Powell (DFP)* [26] formula was proposed as a method to approximate the Hessian matrix, however, this method may not be numerically stable in high-dimensional

hyperparameter searches. Rather than solving a linear system of equations as in the *DFP* the *Broyden-Fletcher-Goldfarb-Shanno (BFGS)* method [27] was proposed to obtain the search direction using matrix multiplication. The result of this is that the *BFGS* is an iterative method which is more numerically stable than the *DFP*, and has very effective “self-correcting properties” [28] and, thus, performs better for solving unconstrained nonlinear optimization problems such as maximizing the log marginal likelihood presented here.

In order to improve the stability and accuracy of the convergence we need to provide the gradient function of each element of the log marginal likelihood. Accordingly, we derive the gradient functions of $\hat{r}, \hat{\omega}, \hat{\sigma}, \hat{\gamma}$ respectively, as,

$$\begin{aligned}
\frac{\partial \ell(r, \omega, \sigma, \gamma)}{\partial r} &= \sum_{i=1}^K \sum_{j=1}^{n_i} \Psi \left[r + \sigma g^{-1}(x_i^T \omega, \gamma) \right] + \Psi(r + y_{ij}) - \Psi(r + y_{ij} + \sigma) - \Psi(r); \\
\frac{\partial \ell(r, \omega, \sigma, \gamma)}{\partial \omega_p} &= \sigma \sum_{i=1}^K \sum_{j=1}^{n_i} \frac{\partial g^{-1}(x_i^T \omega, \gamma)}{\partial \omega_p} (A_i - B_{ij}); \\
\frac{\partial \ell(r, \omega, \sigma, \gamma)}{\partial \sigma} &= \sum_{i=1}^K \sum_{j=1}^{n_i} \left\{ A_i g^{-1}(x_i^T \omega, \gamma) + B_{ij} \left[1 - g^{-1}(x_i^T \omega, \gamma) \right] + C_{ij} \right\}; \\
\frac{\partial \ell(r, \omega, \sigma, \gamma)}{\partial \gamma} &= \sum_{i=1}^K \sum_{j=1}^{n_i} \frac{\partial g^{-1}(x_i^T \omega, \gamma)}{\partial \gamma} (A_i - B_{ij}).
\end{aligned} \tag{17}$$

For clarity, we have defined A_i, B_{ij}, C_{ij} :

$$\begin{aligned}
A_i &= \Psi \left(r + \sigma g^{-1}(x_i^T \omega, \gamma) \right) - \Psi \left(\sigma g^{-1}(x_i^T \omega, \gamma) \right); \\
B_{ij} &= \Psi \left(y_{ij} + \sigma - \sigma g^{-1}(x_i^T \omega, \gamma) \right) - \Psi \left(\sigma - \sigma g^{-1}(x_i^T \omega, \gamma) \right); \\
C_{ij} &= \Psi(\sigma) - \Psi(r + y_{ij} + \sigma),
\end{aligned} \tag{18}$$

where $\Psi(\cdot)$ is the digamma function, such that $\Psi(x) = \frac{\partial \log \Gamma(x)}{\partial x}$, ω_p is the individual element in the vector space ω with $p = (1, 2, \dots, N)$ and N is the number of neurons within the neural population. The gradients of $g^{-1}(x_i^T \omega, \gamma)$ for ω_p and γ required in Eq. (17) are given as

$$\begin{aligned}
\frac{\partial g^{-1}(x_i^T \omega, \gamma)}{\partial \omega_p} &= x_p e^{x_i^T \omega} \left(\gamma e^{x_i^T \omega} + 1 \right)^{-\frac{1}{\gamma}-1}; \\
\frac{\partial g^{-1}(x_i^T \omega, \gamma)}{\partial \gamma} &= - \left(\gamma e^{x_i^T \omega} + 1 \right)^{-\frac{1}{\gamma}} \left[\frac{\log \left(\gamma e^{x_i^T \omega} + 1 \right)}{\gamma^2} - \frac{e^{x_i^T \omega}}{\gamma (\gamma e^{x_i^T \omega} + 1)} \right].
\end{aligned} \tag{19}$$

Although we provide gradient functions to try to ensure stable and accurate optimization, when randomly selecting initial values, it can still sometimes lead to $\gamma < 0$, which does not guarantee a firing probability within $[0, 1]$. To solve this problem, we also enforce the constraint $\gamma > 0$ using Sequential Quadratic Programming (*SQP*) [29], applying both SQP and the Quasi-Newton method at each updating step. To form a quadratic program and find a line search direction by minimizing the quadratic subproblem we first set $\zeta = (r, \omega_{1 \times N}, \sigma, \gamma)$, and form a Lagrangian function as

$$L(\zeta, \mathbf{k}) = \ell(\zeta) + \sum_{m=1}^{N+3} k_m \zeta_m, \quad (20)$$

where, $\mathbf{k} = \{k_1, k_2, \dots, k_{N+3}\}$ is the Lagrangian coefficient. The total number of parameters to be estimated is $N + 3$, N (the number of neurons involved in the functional network) is equivalent to the length of vector ω and (r, σ, γ) are the other three parameters. We then form the quadratic programming subproblem

$$\min_{\mathbf{d} \in \mathbb{R}^d} \nabla \ell^T(\zeta) \mathbf{d} + \frac{1}{2} \mathbf{d}^T \nabla^2 L(\zeta, \mathbf{k}) \mathbf{d}, \quad (21)$$

$$\nabla \zeta_m^T \mathbf{d} + \zeta_m \leq 0, \quad m = 1, \dots, N + 3. \quad (22)$$

These quadratic problems are solved using an active-set algorithm [30]. At each iteration step t , we find the linear search direction \mathbf{d} , and then use a line search procedure to find the step length parameter, Δ , which achieves a sufficient decrease in the merit function. We update the group parameters ζ until converged as below

$$\zeta(t+1) = \zeta(t) + \Delta_t \cdot \mathbf{d}(t). \quad (23)$$

Algorithm 1 summarizes the derivation of the **SODS** estimator and the mean spike count, showing the steps required for empirical Bayes inference on our hierarchical model.

3.2 The Role of the Hyperparameters

The parameters in the **SODS** estimator each play different roles in explaining the dataset. In this section we discuss each of them in turn, and try to present rules to tune the initial values used in the optimization procedure, in order to provide good estimates.

Algorithm 1 Empirical Bayes inference on a hierarchical model.

Input: $\mathbf{x}_i = [x_{i1}, x_{i2}, \dots, x_{iN}]$ and y_{ij} ($i = 1, \dots, K$ and $j = 1, \dots, n_i$).

Output: $\mathbb{E}(\theta_i | \hat{r}, \hat{\omega}, \hat{\sigma}, \hat{\gamma})$ and $\mathbb{E}[Y_{ij} | \theta_i]$.

θ_i is the underlying firing probability of target neuron.

- 1: Initialize $r, \omega, \sigma, \gamma$.
- 2: Form the linear regressors $\eta_i = \mathbf{x}_i^T \omega$.
- 3: Form the extended link function $\eta_i = g(u_i, \gamma)$.
- 4: Derive the prior information as $u_i = g^{-1}(\mathbf{x}_i^T \omega, \gamma)$.
- 5: Construct the Beta Negative Binomial log marginal likelihood function

$$\begin{aligned} \ell(r, \omega, \sigma, \gamma) \propto & \sum_{i=1}^K \sum_{j=1}^{n_i} \left[\log \Gamma(r + \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)) \right. \\ & + \log \Gamma(y_{ij} + \sigma - \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)) + \log \Gamma(r + y_{ij}) \\ & + \log \Gamma(\sigma) - \log \Gamma(r + y_{ij} + \sigma) - \log \Gamma(r) \\ & \left. - \log \Gamma(\sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)) - \log \Gamma(\sigma - \sigma g^{-1}(\mathbf{x}_i^T \omega, \gamma)) \right]. \end{aligned}$$

- 6: Derive the gradient function of each of the parameters

$$\nabla \ell(r, \omega, \sigma, \gamma) = \left(\frac{\partial \ell}{\partial r}, \frac{\partial \ell}{\partial \omega_p}, \frac{\partial \ell}{\partial \sigma}, \frac{\partial \ell}{\partial \gamma} \right).$$

- 7: Apply numerical maximization method *BFGS* to $\ell(r, \omega, \sigma, \gamma)$
- 8: Find the maximum log marginal likelihood estimations of $\hat{r}, \hat{\omega}, \hat{\sigma}, \hat{\gamma}$, calculating the empirical Bayes estimation of firing probability $\mathbb{E}(\theta_i | \hat{r}, \hat{\omega}, \hat{\sigma}, \hat{\gamma})$ as

$$\theta_i^{SODS} = \mathbb{E}(\theta_i | \hat{r}, \hat{\omega}, \hat{\sigma}, \hat{\gamma}) = \frac{n_i \hat{r} + \hat{\sigma} g^{-1}(\mathbf{x}_i^T \hat{\omega}, \hat{\gamma})}{n_i \hat{r} + n_i \bar{y}_i + \hat{\sigma}}.$$

- 9: Obtain the mean spike count based on Eq. (3):

$$\mathbb{E}[Y_{ij} | \theta_i] = \hat{r} \left(\frac{n_i \bar{y}_i + \hat{\sigma} - \hat{\sigma} g^{-1}(\mathbf{x}_i^T \hat{\omega}, \hat{\gamma})}{n_i \hat{r} + \hat{\sigma} g^{-1}(\mathbf{x}_i^T \hat{\omega}, \hat{\gamma})} \right).$$

- r is the number of failures and determines the NB response. In real situations, the actual firing rate of the underlying neural population may not be very high, such as in the hippocampal formation. In this case in order to get a reasonable mean spike count, we should make sure the initial value of r is small, as this helps the spike count observations match the low firing rates. Accordingly, if we believe a brain area has a high firing rate, such as in the motor cortex, we can tune r to be larger initially. For simplicity, we have set r to be the same across all time steps, which means for the whole data length $r_i = r, i = 1, 2, \dots, K$. In the right panel of Fig. 2, we keep the same probability parameter, and show the influence of different values of r on the spike count data. The results indicate that for the NB distribution, larger values of r give larger spike count observations.
- ω is a vector of weights capturing the directed effect of existing neurons on the target neuron. This vector can also include other factors such as the spiking history of the target neuron or external stimuli e.g. if we have prior knowledge, for instance the pixels of an image shown to excite the retinal neurons. These weights can be positive or negative, which can be explained as neurons having either an excitatory or inhibitory effect on the target neuron. In section 4.1, using simulated data, we test the ability of the estimator to capture these excitatory and inhibitory effects of functional connectivity. Without loss of generality, the initial value of each element in ω is randomly chosen from $(-1, 1)$.
- σ is the degree of freedom of the Beta distribution. From Eq. (11), we can see our proposed estimator is the weighted combination of the observed data $\left(\theta_i^{ob} = \frac{n_i r}{n_i r + n_i \bar{y}_i}\right)$ and traditional GLM estimation $\left(\theta_i^{GLM} = g^{-1}(x_i^T \omega, \gamma)\right)$. The weights of each component are $\pi_i = \frac{n_i r + n_i \bar{y}_i}{n_i r + n_i \bar{y}_i + \sigma}$ and $1 - \pi_i = \frac{\sigma}{n_i r + n_i \bar{y}_i + \sigma}$. Thus, if σ is large, the proposed method is close to the GLM of the prior mean; when it is smaller, the estimator is approaching the observed data. The initial value of σ is determined by the number of trials n_i , such that, $\sigma \propto n_i$. This means the more trials we have, the more information we can have about the mean spike counts and the more confidence we can have in the observed data.
- γ selects the best fit link function for our dataset. When $\gamma = 1$, it is the commonly used logit function, when γ is approaching 0, it becomes the complementary log-log link function. Regularly, GLMs choose the link function by specifying a parametric link, our work,

however, determines the unknown parameter by log marginal likelihood. Learning from the dataset itself allows our approach to automatically select a suitable link function. The initial value of γ is determined so as to result in relatively low firing rate, which has empirically been shown to give good performance for mean spike count estimation.

4 Results

In this section, we test our method on both simulated and experimental data. The simulated data is generated via the process outlined in Fig. 1, and the experimental datasets are retinal data taken from four different experiments, each with different numbers of neurons. This data is used to validate the stability and fitness of our model for real spike count data.

4.1 Simulated Data

Our inference method was first tested on data generated from the model shown in the right panel of Fig. 1. The spike-train data is generated from a Beta Negative Binomial model with network sizes 10, 20 and 50. The range of configurations of the parameters tested for the simulated data are shown in the Table 2, where N_s is the number of trials, N the number of neurons in our simulation, and K represents the data length (bins) of each time series in the training dataset. Depending on the combination of these parameters we can provide different types of observed spike count data. Using the simulated dataset, we conducted three tests:

- Interaction estimation. We randomly assigned excitatory or inhibitory weights to the neural population. Our goal was to identify whether we could recover the weights of the interactions.
- Performance of the **SODS** estimator. In the simulation process, we have the underlying truth regarding the mean spike count. We tested the performance of the **SODS** estimator by calculating MSE between the estimation of the mean spike count and the truth.
- Comparison with other modeling methods. Poisson and NB GLMs can also be used to model spike count data, hence, we compared our proposed estimator with each of them by predicting the spike count data of held-out datasets after training each model individually.

Table 2: Configurations of the parameters used for the simulated data.

N_s	N	K	ω	r	σ	γ
1, 10, 20, 50, 100	10, 20, 50	100, 500, 1000, 1500, 2000	(-1,+1)	1, 3, 5	50, 100, 200	1, 3, 5, 7

Table 3 shows a comparison of the MSE of the standard Poisson GLM, the NB GLM and the **SODS** model, with different N_s and K , for one example configuration of the other parameters. Each model is trained using a training set of length K and then tested by computing the MSE between the mean spike count estimation and the true value for a testing set of length 50. The results are averaged across 50 randomly selected initializations of the unknown parameters. The link function of the Poisson GLM and the NB GLM was selected as the *probit* function, we have compared the performance of different link functions (results not shown) such as *log*, *logit*, *probit*, *identity*, and *loglog*, and found the *probit* to give the best performance. Comparing the configurations in Table 3 we find, as would be expected, increasing N_s and K gives a better estimation. When the number of training samples are large all three models have similarly good performance on the simulated data, however, when the number of samples decreases the **SODS** outperforms the Poisson and NB GLMs. The left panel of Fig. 3 shows the effect of having a small number of samples on the **SODS** for different data lengths, while the scatter plot in the right panel of Fig. 3 provides us with a clear view of the comparison between true and estimated mean spike count.

Next, we tested the accuracy of the estimation of the weights describing the directed effect of the neuron population on the target neuron. We explored the effect of different data lengths on the *BFGS* method which was applied to maximize the marginal log likelihood. The results shown in Fig. 4 are taken for several different combinations of the parameter configurations in Table 2, as we can see: (1) the relative standard error is large when the actual weights are close to 0; and (2) 1,000 bins is sufficient to provide accurate and efficient weight estimation in simulation.

4.2 Experimental Data

The experimental data used here is taken from multi-unit recordings of retinal ganglion cells from the ret-1 database [31], curated at CRCNS.org. We tested 4 datasets containing 37, 26, 15, and 14 neurons, respectively. In order to identify the functional dependences among the neural population,

Table 3: Comparison of the Poisson GLM, NB GLM and SODS model for $N = 20$, $r = 5$, $\sigma = 50$, $\gamma = 7$.

$N_s = 10, K = 500$	MSE	$N_s = 50, K = 500$	MSE	$N_s = 100, K = 500$	MSE
Poisson GLM	0.497	Poisson GLM	0.292	Poisson GLM	0.058
NB GLM	0.497	NB GLM	0.292	NB GLM	0.056
SODS Model	0.383	SODS Model	0.204	SODS Model	0.048
$N_s = 10, K = 1000$	MSE	$N_s = 50, K = 1000$	MSE	$N_s = 100, K = 1000$	MSE
Poisson GLM	0.452	Poisson GLM	0.062	Poisson GLM	0.055
NB GLM	0.450	NB GLM	0.062	NB GLM	0.049
SODS Model	0.255	SODS Model	0.052	SODS Model	0.047
$N_s = 10, K = 2000$	MSE	$N_s = 50, K = 2000$	MSE	$N_s = 100, K = 2000$	MSE
Poisson GLM	0.402	Poisson GLM	0.055	Poisson GLM	0.032
NB GLM	0.380	NB GLM	0.050	NB GLM	0.032
SODS Model	0.190	SODS Model	0.047	SODS Model	0.032




Figure 3: Goodness-of-fit for simulated data - *left panel*: the mean square error across data lengths for different numbers of samples at each time step; *right panel*: scatter plot of estimated mean spike count and true mean spike count value with data length 2000.




Figure 4: Scatter plots of estimated weights of 20 neurons ($K = 100, 500, 1000, 2000$). Different colors indicate different neurons. The size of the ball shows the relative standard error (standard error/mean value). The initial values of the weights to be optimized was selected randomly from -1 to 1 and simulated 20 times.

the spike counts were binned into 16.7 ms bins and regressed with the spike counts from the other recorded neurons. From each experimental dataset we created 5 random training and testing splits with a total data length of 2000 bins, 4 splits were used for training and 1 split data were used for testing with 5-fold cross validation performed. The training dataset, consisting of 1600 bins, was used to estimate the unknown parameters of our model and the testing data set of 400 bins, was given as the input regressors to estimate the mean spike count. We then used the **SODS** model, with estimated mean spike count data, to compute the log-likelihoods of the held-out test datasets versus both a standard Poisson GLM and the NB GLM.

Figures 5 and 6 show the percentage increase in the log-likelihood of the **SODS** model over the Poisson GLM and NB GLM for each of the datasets. If we denote ℓ_{SODS} , ℓ_{NB} , ℓ_{Poi} as the predictive log-likelihoods of each method. Then the percentage log-likelihood increase is calculated by $\frac{\ell_{SODS} - \ell_{NB}}{|\ell_{NB}|} \times 100\%$ and $\frac{\ell_{SODS} - \ell_{Poi}}{|\ell_{Poi}|} \times 100\%$. In Figures 5 and 6 a positive value indicates an improvement in the predictive log-likelihood for the held-out dataset of the **SODS** compared to the comparative method and conversely a negative value indicates a decrease in predictive log-likelihood. The majority of results (between 62% and 100% of neurons in each dataset) present higher prediction log-likelihoods for the test data when using the **SODS** model. Notably, the improvement offered by the **SODS** model is larger when compared with the Poisson GLM than when compared with the NB GLM, indicating that the datasets analyzed do indeed have underlying over-dispersed structures. However, when compared with the NB GLM, the **SODS** still offered an increase in performance for a similar number of recorded neurons as for the Poisson GLM. This is despite the NB GLM having the same assumption on the observed spike count data as our **SODS** model.

The **SODS** method outperforms the NB GLM, in particular, when the data length is short. To demonstrate the performance of the different methods under short data length conditions, we analyzed the MSE, averaged across all 14 neurons in dataset #62871, for increasing lengths of the training dataset (the length of the testing set was fixed to 500). The results shown in Table 4, indicate that with a short data length, the **SODS** outperforms the other two models for experimental data. However, the more data provided to the methods the more similar their performance.

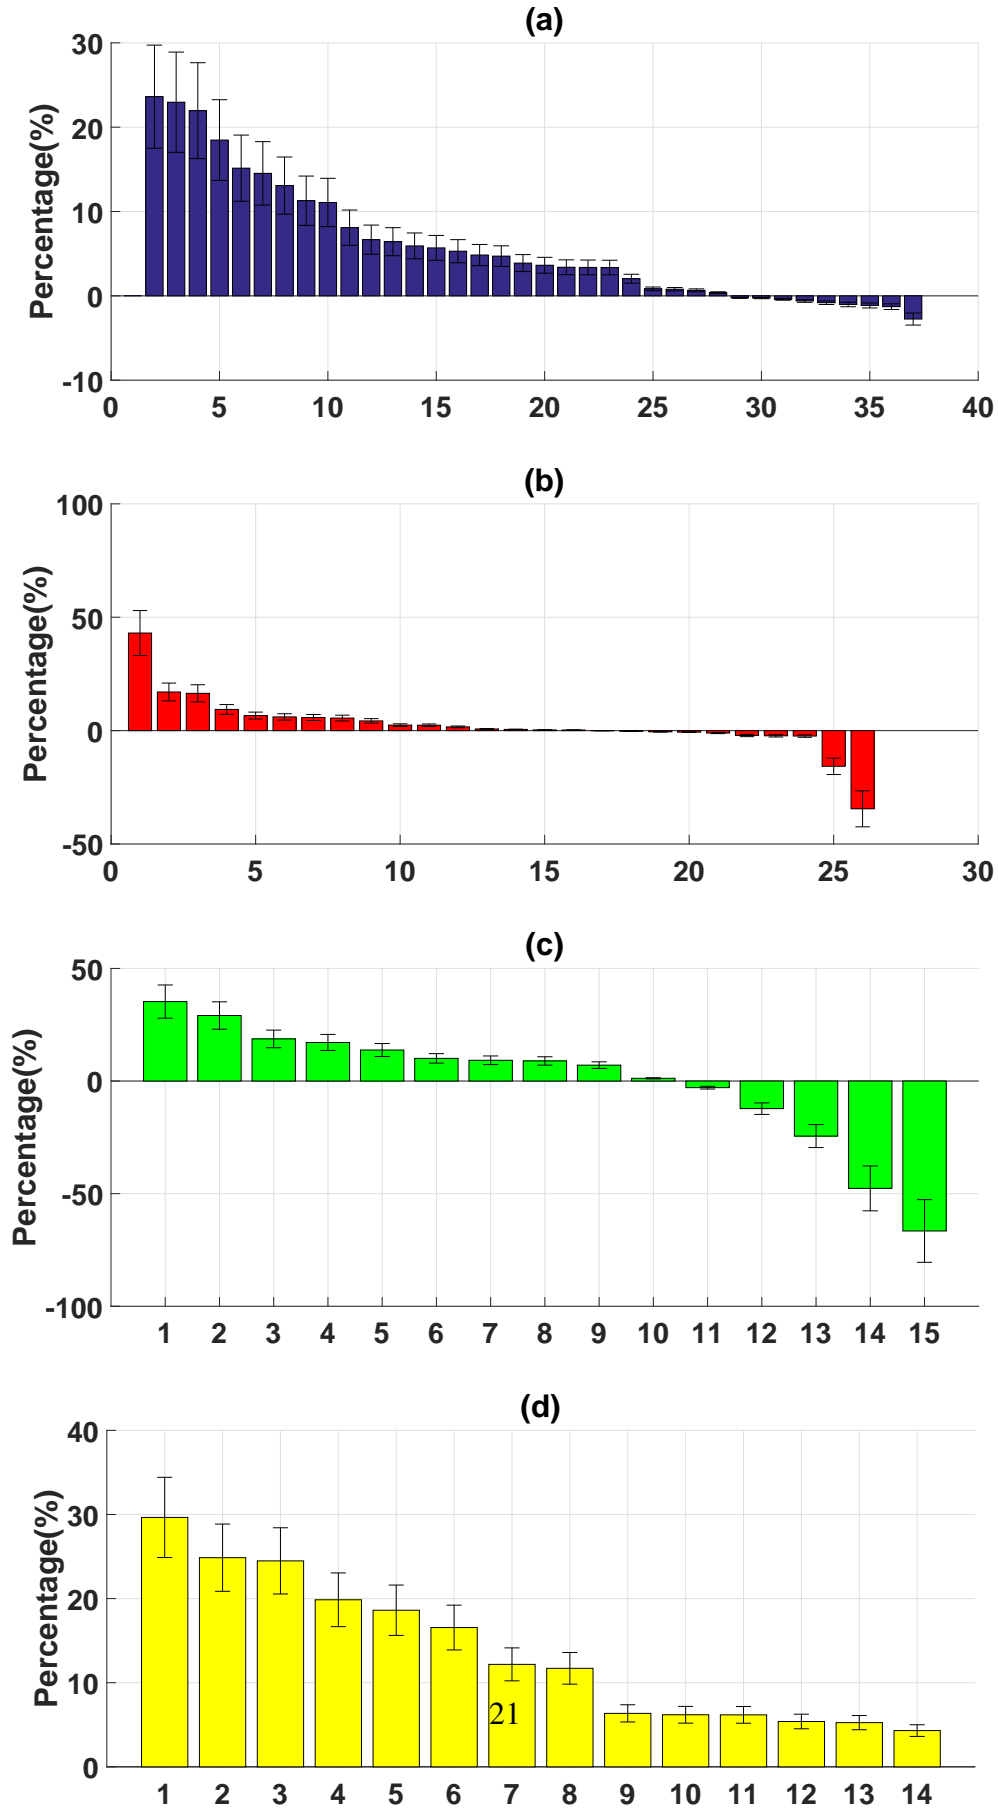


Figure 5: SODS compared with NB GLM: error bar of the percentage increase in the held-out log likelihood of all neurons in dataset **(a)** #62413 (37 neurons); **(b)** #62423 (26 neurons); **(c)** #62814 (15 neurons); **(d)** #62871 (14 neurons).




Figure 6: SODS compared with Poisson GLM: error bar of the percentage in the held-out log likelihood of all neurons in dataset (a) #62413 (37 neurons); (b) #62423 (26 neurons); (c) #62814 (15 neurons); (d) #62871 (14 neurons).

Table 4: MSE averaged across all neurons in dataset #62871.

Methods \ Data length	2,000	3,000	4,000	6,000
Poisson GLM	0.280	0.272	0.260	0.255
NB GLM	0.275	0.267	0.260	0.256
SODS model	0.265	0.261	0.260	0.256

5 Discussion

The neural computation community has shown great interest in multivariate regression methods which can provide a clear view of the nature of relationships, using both linear and nonlinear effects. [32] developed a fully Bayes inference method for NB responses which yields shrinkage for all of the hyperparameters. However, applying fully Bayes approaches to hierarchical models is computationally intensive. As an alternative, empirical Bayes can provide a bias-variance trade-off which can achieve a small mean square error at a lower computational cost. To estimate the unknown parameters of the model, [4] has successfully used maximum likelihood estimation, but when the data length is small, the estimation becomes biased. The **SODS** estimator developed here, to find the latent interactions between neural populations, combines both of the above methods. It has the benefit of providing a shrinkage estimator for NB responses, whilst not needing the intensive computation of fully Bayes inference.

In the work presented in this letter, we take advantage of the beneficial properties of both GLMs and empirical Bayes inference to propose the **SODS** estimator. We used the NB distribution to model the spike count process of each neuron. The NB distribution was selected as it allows for over-dispersion of the spike count using a dispersion parameter superior to that of the standard Poisson model. At the same time, the Beta distribution is employed, as the prior information for the probability parameter in the NB distribution, to allow for better posterior distribution derivation. We utilize a flexible link function in order to model the prior mean using regressors. By using the recorded data from other neurons as the covariates, we can then infer the functional weights among the neural population. Unlike fully Bayes inference which utilizes hyperpriors, we instead estimate the hyperparameters by maximizing the marginal likelihood. The result is that the

proposed **SODS** estimator is a shrinkage estimator and the weights we estimate can be viewed as the hidden functional dependences. By taking the neurons as nodes in our functional network, and their spike-train data as the observations, our empirical Bayes inference method can be used to identify the neural interactions, including excitatory and inhibitory behaviors. The **SODS** estimator also provides a valid method to estimate the underlying firing probability of each neuron and estimate corresponding mean spike count.

We have validated our method using both simulated data and experimental retinal neuron data. We compared the proposed estimator against GLM estimators, using both Poisson and NB regressions. By using intensive simulations we have shown, that on our simulated system, the **SODS** outperforms the Poisson and NB GLMs. The proposed approach implements a flexible link function, which unlike the Poisson and NB GLMs, means it can select the best link for each dataset. From the simulation data we found that by efficiently maximizing the marginal likelihood, we can accurately estimate the model parameters. For the experimental data, when compared with two of the most widely used regression methods: Poisson and NB regressions, there was substantial improvement in the log likelihood of the held-out datasets. Whilst the results presented here are promising, going forward, we are interested extending our model, for instance with the incorporation of Hebbian learning rules which would allow us to account for time-varying weights. Applying prior knowledge regarding network structure such as random, small world or scale-free networks could also be a promising avenue for future research.

6 Conclusion

Recently, NB models have been used in systems neuroscience due to their suitability for over-dispersed data [20, 25]. However, difficulties arise when using Bayesian methods in conjunction with hierarchical NB modeling, therefore, our work addresses three associated contributions. First, we have proposed a shrinkage **SODS** estimator with a lower MSE than alternative GLM methods, and have validated it using both simulated and experimental data. Second, our proposed method is ideally suited to short over-dispersed spike count data. While, at the same time, our approach avoids the high computational costs often associated with using fully Bayes methods by employing empirical Bayes inference. Third, the **SODS** method has better performance, when compared

with both Poisson regression and NB regression methods, in terms of: (1) identifying interactions among neural populations; and (2) estimating mean spike counts. Our approach can also easily be extended to take more covariates into consideration.

Acknowledgments

The work described in this letter was fully supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. CityU110813].

References

- [1] H.-J. Park and K. Friston, “Structural and functional brain networks: from connections to cognition,” *Science*, vol. 342, no. 6158, p. 1238411, 2013.
- [2] M. Okatan, M. A. Wilson, and E. N. Brown, “Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity,” *Neural Computation*, vol. 17, no. 9, pp. 1927–1961, 2005.
- [3] J. W. Pillow, L. Paninski, and E. P. Simoncelli, “Maximum likelihood estimation of a stochastic integrate-and-fire neural model,” in *Advances in Neural Information Processing Systems*, pp. 1311–1318, 2003.
- [4] L. Paninski, “Maximum likelihood estimation of cascade point-process neural encoding models,” *Network: Computation in Neural Systems*, vol. 15, no. 4, pp. 243–262, 2004.
- [5] W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown, “A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects,” *Journal of Neurophysiology*, vol. 93, no. 2, pp. 1074–1089, 2005.
- [6] J. W. Pillow, J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. Chichilnisky, and E. P. Simoncelli, “Spatio-temporal correlations and visual signalling in a complete neuronal population,” *Nature*, vol. 454, no. 7207, pp. 995–999, 2008.

- [7] Q. She, W. K. So, and R. H. Chan, “Reconstruction of neural network topology using spike train data: Small-world features of hippocampal network,” in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 2506–2509, IEEE, 2015.
- [8] C. M. Bishop, M. Svensén, and C. K. Williams, “GTM: The generative topographic mapping,” *Neural Computation*, vol. 10, no. 1, pp. 215–234, 1998.
- [9] N. D. Lawrence, “Gaussian process latent variable models for visualisation of high dimensional data,” in *Advances in Neural Information Processing Systems*, no. 3, pp. 329–336, 2004.
- [10] N. Lawrence, “Probabilistic non-linear principal component analysis with gaussian process latent variable models,” *The Journal of Machine Learning Research*, vol. 6, pp. 1783–1816, 2005.
- [11] K. C. Lakshmanan, P. T. Sadtler, E. C. Tyler-Kabara, A. P. Batista, and M. Y. Byron, “Extracting low-dimensional latent structure from time series in the presence of delays,” *Neural Computation*, vol. 27, no. 9, pp. 1825–1856, 2015.
- [12] J. H. Macke, L. Buesing, J. P. Cunningham, M. Y. Byron, K. V. Shenoy, and M. Sahani, “Empirical models of spiking in neural populations,” in *Advances in Neural Information Processing Systems*, pp. 1350–1358, 2011.
- [13] M. M. Churchland, M. Y. Byron, J. P. Cunningham, L. P. Sugrue, M. R. Cohen, Corrado, *et al.*, “Stimulus onset quenches neural variability: a widespread cortical phenomenon,” *Nature Neuroscience*, vol. 13, no. 3, pp. 369–378, 2010.
- [14] J. F. Lawless, “Negative binomial and mixed poisson regression,” *Canadian Journal of Statistics*, vol. 15, no. 3, pp. 209–225, 1987.
- [15] Z. Chen, D. F. Putrino, S. Ghosh, R. Barbieri, and E. N. Brown, “Statistical inference for assessing functional connectivity of neuronal ensembles with sparse spiking data,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 19, no. 2, pp. 121–135, 2011.

- [16] G. M. Cordeiro and P. McCullagh, “Bias correction in generalized linear models,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 53, no. 3, pp. 629–643, 1991.
- [17] M. D. Mauk and D. V. Buonomano, “The neural basis of temporal processing,” *Annual Review of Neuroscience*, vol. 27, pp. 307–340, 2004.
- [18] N. Bertschinger and T. Natschlager, “Real-time computation at the edge of chaos in recurrent neural networks,” *Neural Computation*, vol. 16, no. 7, pp. 1413–1436, 2004.
- [19] Q. She, G. Chen, and R. H. Chan, “Evaluating the small-world-ness of a sampled network: Functional connectivity of entorhinal-hippocampal circuitry,” *Scientific reports*, vol. 6, 2016.
- [20] J. Scott and J. W. Pillow, “Fully bayesian inference for neural models with negative-binomial spiking,” in *Advances in Neural Information Processing Systems*, pp. 1898–1906, 2012.
- [21] M. C. Wiener and B. J. Richmond, “Decoding spike trains instant by instant using order statistics and the mixture-of-poissons model,” *The Journal of Neuroscience*, vol. 23, no. 6, pp. 2394–2406, 2003.
- [22] M. R. DeWeese, M. Wehr, and A. M. Zador, “Binary spiking in auditory cortex,” *The Journal of Neuroscience*, vol. 23, no. 21, pp. 7940–7949, 2003.
- [23] Y. Gao, L. Busing, K. V. Shenoy, and J. P. Cunningham, “High-dimensional neural spike train analysis with generalized count linear dynamical systems,” in *Advances in Neural Information Processing Systems*, pp. 2044–2052, 2015.
- [24] I. H. Stevenson, “Flexible models for spike count data with both over-and under-dispersion,” *Journal of Computational Neuroscience*, pp. 1–15, 2016.
- [25] R. L. Goris, J. A. Movshon, and E. P. Simoncelli, “Partitioning neuronal variability,” *Nature Neuroscience*, vol. 17, no. 6, pp. 858–865, 2014.
- [26] D. Goldfarb, “A family of variable-metric methods derived by variational means,” *Mathematics of Computation*, vol. 24, no. 109, pp. 23–26, 1970.
- [27] D. F. Shanno, “On broyden-fletcher-goldfarb-shanno method,” *Journal of Optimization Theory and Applications*, vol. 46, no. 1, pp. 87–94, 1985.

- [28] H. W. Kuhn, *Nonlinear programming: a historical view*, pp. 393–414. Springer, 2014.
- [29] Y. Taniai and J. Nishii, “Optimality of upper-arm reaching trajectories based on the expected value of the metabolic energy cost,” *Neural Computation*, vol. 27, no. 8, pp. 1721–1737, 2015.
- [30] N. Gillis and F. Glineur, “Accelerated multiplicative updates and hierarchical ALS algorithms for nonnegative matrix factorization,” *Neural Computation*, vol. 24, no. 4, pp. 1085–1105, 2012.
- [31] J. L. Lefebvre, Y. Zhang, M. Meister, X. Wang, and J. R. Sanes, “ γ -protocadherins regulate neuronal survival but are dispensable for circuit formation in retina,” *Development*, vol. 135, no. 24, pp. 4141–4151, 2008.
- [32] S. W. Linderman and R. P. Adams, “Discovering latent network structure in point process data,” in *Proc. International Conference on Machine Learning*, pp. 1413–1421, 2014.